


Focusing On IBM Host-Based Enterprise-Wide Computing

ENTERPRISE SYSTEMS™

J · O · U · R · N · A · L

December 1995



Data Mining:
Techniques To Gain
Insight Into Your Data


ESJ CONTENTS

December 1995 • Volume 10, Number 13

Learn about the internal workings of VSAM on page 72.



High-performance, low-latency switching solutions on page 48.



See page 18 for details on how to use data mining to extract information — the “gold” — from the larger data “ore.”
Cover illustration by Robert Kimmerle.

OPERATING SYSTEM PLATFORMS

- CICS/ESA 4.1 Pioneers *By Joseph Larson* 66
The Anatomy Of VSAM CI And CA Splits *By Jim Lewandowski* 72

CLIENT/SERVER ENTERPRISE SOLUTIONS

- In Praise Of PL/I *By Lou Marco* 32
ISPF 4.2: The Momentum Continues *By Jon E. Pearkins* 38

DATABASE MANAGEMENT

- Data Mining: Techniques To Gain Insight Into Your Data
By Emil T. Cipolla 18
Gathering DB2 Information In Cyberspace *By Craig S. Mullins* 26

ELECTRONIC DOCUMENT MANAGEMENT

- Chep USA Teams VM And OS/2: Automates 1000 Faxes
By Patrick Sharkey 70

INTERNETWORKING SECTION

- Interview With Gary Hegna Of OpenConnect Systems 44
Protect Your Investments In Token-Ring Networks
By Marthin deBeer 48
Reliable Gigabyte Networking Solutions: Disaster Recovery Backup
And Client/Server Information Transfers *By Danny Quam* 52

Company Profiles:

- | | |
|--|--------------------------------------|
| <i>Barr Systems, Inc.</i> 55 | <i>McGill Systems, Inc.</i> 59 |
| <i>Computerm</i> 56 | <i>NetManage, Inc.</i> 60 |
| <i>Data Interface Systems Corp.</i> 57 | <i>OpenConnect Systems</i> 61 |
| <i>Levi, Ray & Shoup, Inc.</i> 58 | <i>Renex Corp.</i> 62 |

DISTRIBUTED COMPUTING SOLUTIONS

- The Nature Of The Beast: Understanding What It Takes To Manage
Networked Computing Systems *By Richard Brooks* 10

FIRST IMPRESSION

- Open PL/I (*Liant Software Corp.*) 8

DEPARTMENTS

- Publisher 4
Inside IBM 6
Advertiser Index 64
Humor 78
Viewpoint 80

Chep USA Teams VM And OS/2

Distributing manufactured goods on wooden pallets in this country carries the heavy load of \$2.5 billion in waste. Of the 100 million pallets in circulation in the grocery industry, 66 million are thrown away every year at a cost of \$700 million. The abundance of pallets in poor condition wastes \$2 billion annually because of their contribution to product damage, according to a recent study by the U. S. grocery industry.

The billions of wasted dollars presented a golden opportunity for any company capable of efficiently managing a manufacturer's pallets. Throughout the world, Chep has taken this opportunity, helping companies minimize their expenses in the distribution of pallets. In 1990, Chep came to the United States to begin similar operations at the request of Procter & Gamble, which was a client of Chep's in Europe.

Keeping Track Of 12 Million Pallets

Currently, Chep USA (Park Ridge, NJ) has approximately 12 million pallets in circulation throughout North America. According to William Homa, vice president of Information Technology, keeping track of this inventory requires that Chep make 5000 communications daily — notifying manufacturers and distributors of how many pallets they will receive and when, alerting Chep's 165 depots of the movement of pallets, letting out bids for trucking movements and so on. The company tracks pallets by transaction, just as a bank tracks money. Much of its man-



Automates

1000

Daily Faxes

By Patrick Sharkey

agement effort is in relocating pallets, since most of the manufacturing is in the center of the country while most of the consumption is on both coasts. To "re-balance the pool," Homa says, they relocate more than 750,000 pallets, or 1300 truckloads, each week.

"When we started U.S. operations in September 1990, we had to be nationwide from day one," Homa points out. "We developed an infrastructure of people and systems before we started." To keep track of pallets, they initially used an IBM 4381 mainframe running VM. Pallet management was the mainframe's sole application. Homa says they used VM because they adapted code developed by their French operations for this operating system. When an American company runs VM, it is usually with MVS as a guest operating system. In Europe, however, it is more common to run a business application natively on VM.

The advantage of Chep's code and choice of operating system was that as it upgraded its mainframe, the company did not have to change one line of code, since it was scalable. Recently, Chep upgraded to a Hitachi GX 6215 mainframe, which runs from Michigan and is connected to the company's Park Ridge office by a WAN. Last year, Chep's operations in Canada, Mexico and Chile were all connected to the Hitachi by a leased-line network. Chep also uses IBM AS/400 minicomputers as remote communications servers and to run accounting applications.

Communicating By EDI And Faxes

Fortunately, the majority of Chep's transactions are received and processed by Electronic Data Interchange (EDI), in industry-standard ANSI X.12 format through the gateway of a third-party vendor. While most of Chep's customers communicate with the company during overnight EDI batch runs and many of Chep's carriers are linked using realtime transaction sets, the company was unable to increase its EDI penetration above 80 percent. The reason was that a substantial number of its business partners, such as trucking firms and depots, did not have access to EDI.

To reach these partners, Chep had to use everyman's communications standard: the fax machine. In spite of its EDI

success, Chep's volume in faxes has been about 1000 daily, most of them sent during five peak hours of the workday. That works out to four faxes each minute.

To handle this load, Homa explains, Chep had employees whose primary job was to send out faxes. They printed information from the mainframe, took the document to the nearby fax and frequently were held up by the congestion. Occasionally, documents that were already in the fax machine, on hold while the machine waited to redial the number after encountering a busy signal, were removed and replaced after the second document was sent.

As Chep's business continued to grow, the number of faxes was not going to decrease. It had to find a way to automate fax transmissions. Chep tried a DOS-based system that claimed to automate the transmission of faxes from the mainframe. Its benefits, however, were limited. Homa says that because of a weak link in the VM application, the sender had to manually print a page to find out if the fax was successfully sent.

More important, the DOS foundation rendered the product unreliable. DOS was not robust enough to handle the many simultaneous operations involved in the process. The fax server downloaded documents from the mainframe, scanned the documents into a graphic format, dialed the fax number, sent out the fax and determined if it was successful. During all of this, the program had to maintain a handshake with the mainframe. It was asking DOS to do too much. Another problem was that the vendor did not have experience working with VM, so it was at a loss to make quick fixes to the VM interface. The DOS product missed the key requirement to have the fax application feed back to the VM application the status of a particular fax — a vital component since Chep's typical user is 1000 miles away from the fax PC.

While searching for another solution, Chep found an OS/2-based fax server made by a vendor with VM experience (Software Pursuits, Alameda, CA). Chep knew OS/2 would be more reliable than DOS, since it is a full, preemptive multitasking system with none of the memory limitations of DOS or Windows to cause errors in the execution of the program.

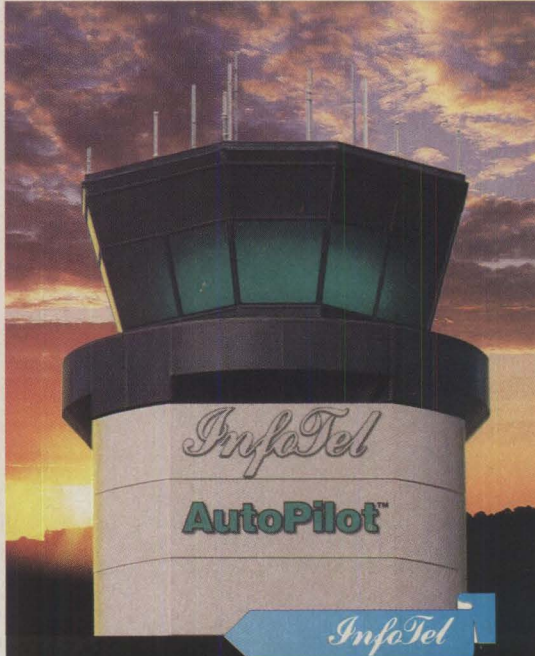
See Chep USA on page 76

Take complete control of DB2® data migration

INFORIM2 enables users to easily extract and migrate referentially and relationally intact subsets of DB2 data which span any number of DB2 tables. This saves a tremendous amount of time as the user has complete control over exact rows of data to be migrated and can complete an entire execution from a single request.

Call InfoTel today at **800-543-1982** and ask about INFORIM2 and our entire suite of simple-to-use mainframe software utilities—**INFOTRACE, INFOPAK, INFOLOAD, INFOSCAN, INFOSCOPE, INFOCOPY, INFOCYPHER, INFOREORG, and INFORECOVERY.**

PUT YOUR DATA CENTER ON AUTOPILOT!



InfoTel
AutoPilot™

InfoTel
Simple Brilliance in Software

15438 N. Florida Ave. Ste204
Tampa, FL 33613
Phone: 800-543-1982 Fax: 813-960-5345
Internet: 76544.2456@CompuServe.com

AutoPilot™

DB2 is a registered trademark of International Business Machines Corp.
CIRCLE #186 on Reader Service Card ▲


serialize them. This may shed light on performance problems for on-line data sets that incur huge amounts of splits in light of the fact that VSAM must serialize all split activity for a data set. In effect, the I/O involved for CI split processing is not the true bottleneck as much as the fact that CI splits are serialized by VSAM.

Also, during any split processing, the CA is locked by exclusive use of the index sequence set CI to defer other requests from its use. In a system with high levels of activity, these accesses to the CA are deferred by VSAM until split processing is complete for the CA.

The time for a CA split — less than one second to three seconds or more — may not seem like a large bottleneck at first glance. For highly active data sets, however, requests coming in for the split CA would have to be deferred. These deferred requests could use up available Placeholder (PLH) string resources and cause other requests to double up on waiting for a PLH. Conversely, even with CI splits occurring well under subsecond, the same situation could occur in a less severe form. Even if a data set is difficult to tune for split activity, a clear picture of quantifying the elapsed time can help to understand the performance implications of splits.

This article has illustrated what occurs during VSAM CI and CA split processing. It should be evident that VSAM was designed with as much thought to throughput and performance as possible. Clearly, CA splits incur, on average, 100 times the I/O overhead of a CI split. This understanding should help provide concrete reasoning during discussions pertaining to tuning VSAM data sets with respect to splits. ●

ABOUT THE AUTHOR



Jim Lewandowski is president of GLH, Inc., an independent consulting company. He has 12 years experience in I/S with expertise in VSAM, software development, MVS systems performance, application performance, REXX/Dialog manager, vendor software product evaluation and catalog information processing via SVC 26. GLH, Inc., 229 W. Lincoln Ave., Barrington, IL 60010, (708) 381-5817.

Like UNIX, OS/2 has never had the 640K main memory barrier that plagues DOS and thus can offer all available memory for an application.

Chep tested the fax server and its interaction with VM and found the application efficiently handled the concurrent processes involved in downloading information and sending faxes. The fax server PC, directly connected to the Hitachi mainframe in Michigan, is a 486-66 Dell PC with 16MB of RAM and a 350MB hard disk. The fax hardware is a four-line Brooktrout (Needham, MA) TR 114 internal fax modem. Along with OS/2 and the Software Pursuits Fax/DSF software, the PC runs IBM Communication Manager/2 for the mainframe link.

Unlike the DOS-based fax product, the Software Pursuits fax server was able to update the mainframe on whether the fax transmission was successful. The vendor worked with Chep to add the capability of notifying the sender if the fax was unsuccessful, a feature that is now a standard part of the product. They also jointly wrote an application to automatically monitor messages sent to the fax administrator and notify the fax originator if the fax failed to transmit within 20 minutes.

Additionally, Chep developed a mechanism to page a systems programmer in New Jersey if faxes began to accumulate in the data center where the fax server was situated or if there was any other problem, such as a lost circuit. When the programmer receives the page, he can log into the system and check the queue. He can also log onto the PC using OS/2 remote control software and operate the computer, rebooting it if necessary. If there is a problem, it can generally be fixed remotely from New Jersey.

In addition to its reliability, Homa says, the OS/2 fax server offered a number of user-friendly features for which the PC platform is known. First of all, the fax pages can be enhanced with a custom forms overlay. In Chep's case, it added the Chep USA logo in the right corner and boxes in the middle of the page for the recipient's name, the subject of the memo and the memo. The sender does not have to call up this form; it is automatically included with the document.

The system also lets users easily send a fax on their own. One of the staff, for example, may want to send a fax to

confirm a conversation or an order. Using the mainframe as a word processor, the person can type a memo, even add mainframe data and send the document as a fax without having to log off the mainframe session. The user simply "prints" it to the fax server. This gives added convenience to the staff because most of them do not have faxes in their immediate areas. Another feature is that the sender can change the page orientation of the document from conventional portrait mode to the landscape orientation of spreadsheets. That is handy when the sender is presenting several columns of figures.

The Savings Of A Reliable Fax Server

The reliability and the ease of use of the Fax/DSF system is saving Chep \$250,000 a year, according to Homa. This savings is mostly in labor and office space. Next year, the savings will be greater as volume continues to increase. The fax server has not replaced all manual effort, since Chep continues to receive paper faxes and has to route them to the appropriate staff. The administrative workload, however, is diminished from what it was previously.

With the success of the fax server, Chep recently purchased a second unit. It will use this unit, which is hooked to its token-ring network in New Jersey, to split the workload between the two PCs, mostly to accommodate the peak demand during the day.

"The reliability of the remote fax server has helped us accomplish our goal of running a 'lights-out,' completely unattended data center," Homa notes. Chep operations in Canada, Mexico, Chile and its offices in New Jersey are all connected by WAN to the data center in Michigan. It is somewhat unusual to have a VM system running unattended 24 hours a day, seven days a week. Chep's runs problem-free.

In the future, a greater percentage of Chep transactions will be EDI. As its business continues to expand, however, its fax volume will still be substantial. Without the fax server, Homa says, they simply could not efficiently handle the high volume of faxes. ●

ABOUT THE AUTHOR

Patrick Sharkey is a freelance writer based in Detroit, MI.